AD-A089 346   SYSTEM DEVELOPMENT CORP MCLEAN VA   F/G 15/5
A COMPARISON OF RISK ASSESSMENT METHODOLOGIES.(U)
JUN 80 H S MORSE   N00173-78-C-0455
UNCLASSIFIED SDC-TM-WD-7999/202/00   NL

1 OF 1
AD A
089346

END
DATE
FILMED
10-80
DTIC

2

AD A089346

LEVEL II

DTIC ELECTE SEP 22 1980 S C D

A COMPARISON OF RISK
ASSESSMENT METHODOLOGIES

JUNE 30, 1980

This document has been approved for public release and sale; its distribution is unlimited.

TM-WD-7999/202/00

DDC FILE COPY

80 9 18 075

A COMPARISON OF RISK ASSESSMENT METHODOLOGIES

BY: H. STEPHEN MORSE

JUNE 30, 1930

TM-WD-7999/202/00

ABSTRACT

This document contains a comparison of five Risk Assessment Methodologies. They are: FIPS 65; the Air Force RAMP; Chapter 6 of the DoA ADP Handbook; the AFIPS Self-Audit Checklist; and the RAM developed by SDC for the Navy. These methodologies are compared on a number of criteria including theoretical soundness, level of effort required, ease of use, completeness of instructions, use and reliability of quanitative metrics, and appropriateness for use by the Department of the Navy.

TABLE OF CONTENTS

Paragraph Page

SECTION 1 - INTRODUCTION

SECTION 2 - COMPARISON OF METHODOLOGIES

SECTION 3 - CONCLUSIONS

1. INTRODUCTION

This document presents a comparison of some currently available ADP Risk Assessment Methodologies. The methodologies which are compared are:

1. Guidelines for Automatic Data Processing Risk Analysis (8/79)
   FIPS PUB 65
   National Bureau of Standards, Washington, D.C.

2. Risk Analysis Management Program (no date)
   Vol. I-III
   AF REGULATION 300-XX
   Department of the Air Force, Washington, D.C.

3. ADP Security Handbook (8/77)
   FIPS Manual Chapter 6
   Department of Agriculture, Washington, D.C.

4. Security: Checklist for Computer Center Self-Audits (79)
   Peter S. Browne
   AFIPS System Review Manual
   AFIPS Press, Arlington, VA

5. Risk Assessment Methodology (7/79)
   TM-WD-7999/001/03
   System Development Corporation, McLean, Virginia

The remainder of this section presents the criteria on which the comparison will be based (paragraph 1.1), and a high-level overview of each of the methodologies (paragraph 1.2). The comparison itself is presented in Section 2, organized by comparison criteria. Section 3 contains conclusions and recommendations.

1.1 COMPARISON CRITERIA

It has been generally recognized for some time that ADP managers require a systematic, quantitative approach to computer security and risk management. The tool which has evolved to meet this need is the Risk Assessment. The underlying idea is quite simple: obtain reliable estimates of the frequency of occurrence of threats to a facility or system, and of the impact (cost) associated with their occurrence. With these estimates, selection of cost-effective countermeasures and operational procedures can be carried out with increased confidence.

To have enunciated this concept, however, is not to have the product in hand, and computer Risk Assessments (like any other imprecise procedure) are plagued with a number of major and minor difficulties. Among them are:

- Because of the technical nature of much of the analysis, a Risk Assessment may require highly skilled (and expensive) manpower.

- Many of the critical parameters which enter into a Risk Assessment are extremely difficult to estimate with precision.

- On the one hand, a quantitative measure (like dollars-per-year) is very useful in justifying additional expenditures for protective countermeasures. On the other hand, many system assets (classified data, loss of personnel, etc.) do not readily lend themselves to such evaluation techniques.

- In many methodologies, it is unlikely that Risk Assessments conducted by different people or on different facilities are comparable.

o It is difficult to build assurance that all of the major threats to and vulnerabilities of a system have been taken into account.

These general areas of difficulty point toward useful comparison and evaluation criteria for Risk Assessments. The extent to which a methodology recognizes and solves these problems is an indication of its technical merit and operational feasibility. The comparison criteria should be broad enough to deal with these issues.

This comparison is being conducted for the Department of the Navy, and the unique mission and ADP requirements of the Navy impose additional constraints and requirements on a Risk Assessment Methodology. The appropriateness for use by Navy personnel, issues of ADP resource availability, the level of effort and expertise, and the reliability of results are all of special importance in the context of Navy ADP operations. Comparison criteria should also reflect these concerns.

A careful review of these considerations and of the Risk Assessments has led to five major areas for comparison. Together, they cover the most significant points of difference, and should give a reliable picture of the relative technical merits of each of the methodologies. This section will conclude with a brief discussion of each comparison criterion.

1. Soundness of underlying theory: Each of the methodologies presents an algorithm for estimating certain parameters and combining those estimates into summary statistics. These include procedures for gathering data, organizing it, estimating parameters, computing, and analyzing the findings. The procedures which are called for, in turn, arise from an underlying model of the risk environment of an ADP system or facility. The validity of this model, and the

extent to which the methodology as presented consistently reflects that model, are pivotal to the overall reliability and utility of the quantitative measures which are produced. Mathematics only reflects the real world if care and insight have been used in constructing the model.

2. Level of manpower required: In many Risk Assessments, the directions for conducting the assessment are somewhat vague and open-ended. Therefore, the question of what level of effort is required to obtain reasonable results is an important one. In particular, a methodology which can produce good results with substantially less effort than alternative approaches has a marked advantage over them. Also, the level of ADP expertise required to perform the Risk Assessment is a significant consideration.

3. Ease of use, Completeness of instructions and Designation of responsibility: A good risk assessment methodology can substantially ease the burdens of performing the assessment and evaluating its results. The mechanisms for this are 1) setting forth complete and detailed instructions for performing the risk assessment; 2) providing detailed guidance on estimation of parameters whenever possible; and 3) clearly designating responsible individuals in the areas of data collection, estimation, support, review and final action. The methodologies will be compared in this regard.

4. Use and reliability of quantitative metrics: A risk assessment is an attempt to quantify the level of risk present at a facility or system. The methodology should produce these measures in a way which is consistent, complete and reliable. A methodology which 1) does not reflect major areas of concern; 2) can be conducted with widely varying results; or 3) fails to quantify the relevant attributes, is, to that extent, deficient.

5. Appropriateness for use by the Department of the Navy: Any feature exhibited by a methodology which seems to be particularly well-suited to the needs and mission of the Navy will be discussed under this heading.

1.2 OVERVIEW OF THE METHODOLOGIES

1.2.1 FIPS PUB 65

This document is the natural successor to FIPS PUB 31 which was for some time the principal reference for Federal ADP managers on risk assessment in their facilities and systems. It incorporates several changes, including the use of orders of magnitude (a la Courtney), a countermeasure selection mechanism, and an appendix which contains an extensive list of ADP system vulnerabilities. The intended audience is Federal ADP managers, but the strategy and techniques are sufficiently general to be potentially of use in a wide variety of settings. The document is very short (41 pages) and can be profitably read in under an hour. The following features merit special attention.

1. Purpose: "The aim of a risk assessment is to help ADP management strike an economic balance between the impact of risks and the cost of protective measures . . . A secondary benefit of a risk assessment is the increased security awareness which will be apparent at all organizational levels ..."

2. Level of effort: "The major resource for a risk analysis is manpower -- highly skilled manpower . . . If meaningful results are expected, management must be willing to commit the resources necessary for accomplishing this undertaking."

3. Management participation: The document stresses the need for management support and participation, including the selection of

a qualified team to perform the risk analysis, setting aside sufficient time to conduct the analysis, review of findings and resolution of potential conflicts.

4. Three staged approach: The methodology calls for three basic activities: A) a preliminary examination; B) the risk analysis proper; and C) the selection of cost-effective countermeasures.

A) In the Preliminary Security Analysis, the team identifies the replacement costs for system assets, the threats (by name) to which the system is subject, and the existing security measures. While these three discrete activities are called for, guidance on actually doing them is very sketchy and incomplete. Some examples of threats and system vulnerabilities are given, but the structuring of these activities and the level of detail is left to the judgement of those conducting the assessment. This is one example of the loosely structured nature of the methodology as presented in FIPS PUB 65. The final activity in step A) is a management review of the preliminary findings.

B) Risk Analysis: The procedure calls for estimation of two quantities: frequency of occurrence of a threat, and impact in dollars when a threat affects some asset. Multiplication of these factors results in the annual loss which can be expected from that threat. Summing over all threat-asset pairs provides the annual loss expectancy (ALE). The ALE is the principal summary statistic resulting from the risk assessment. It is a statistical average over time, and it may take a number of years before actual losses agree with average losses.

While the document explicitly recognizes the need for a structured approach, it offers only marginal guidance on how to actually

implement such structure. A form (Risk Analysis Worksheet) is presented which provides entries for computation of the ALE in the impact areas of modification of data, destruction of data, confidentiality of data (i.e., unauthorized disclosure) and processing availability (denial of service). The ALE is computed by asset, and then summed over all assets to obtain the system-wide ALE. Thus, for each system asset, it is necessary to a) identify each threat which could impact the assets; b) estimate the frequency of occurrence of the threats in each impact category; c) estimate the impact in dollars of that threat in each of the four impact categories; d) multiply these ratings (using orders of magnitude) to obtain the ALE for each asset-threat pair; and e) sum over all such asset-threat pairs to obtain the system-wide ALE. Within the denial of service impact mode, three degrees of impact are suggested (the example uses 2 hours, 24 hours and 72 hours).

Note that the fundamental indexing entity is the asset. For some reason, it is taken for granted that data files will be the most prominent assets. While they certainly do constitute an important category of assets, other assets (i.e., equipment, software, personnel, communications, negotiable output, etc.) should not be neglected.

Another significant feature of the risk analysis is the following: "The effect of currently installed protective measures on undesirable events should not be taken into account at this stage." Thus, the team is asked to hypothesize what the frequency and impact would be if all currently installed countermeasures were removed.

An example of this process is provided.

C) Selection of Safeguards: The emphasis is heavily on procedural and physical safeguards: "Procedural controls, especially when used in combination with physical barriers, produce the highest degree of security for the lowest cost of all forms of protection . . . System security measures should be contemplated only after it has been established that physical and procedural safeguards are insufficient to meet the organization protective requirements."

It is at this stage that in-place protective measures are considered, including replacement costs. By comparing the annual cost of a countermeasure against the reduction in ALE it produces, cost effective protective measures and devices can be selected. Little guidance is given on how to arrive at reasonable prospective safeguards. The procedure calls for each such measure to be matched against each threat (in a matrix), and the reduction of ALE estimated in each instance. An example is provided.

5. Appendix Containing System Vulnerabilities: One of the most useful parts of the document is an eight-page Appendix listing "many undesirable events which can have serious consequences." These are arranged under such headings as Uncontrolled System Access, Procedural Errors, Program Errors and Communications Failure.

An overall technical assessment of this methodology must include the observation that FIPS PUB 65 occupies a unique place in the field of Risk Assessment. The prestige and technical expertise of NBS carries significant influence in the ADP community. As a high-level statement of the purposes and methods of a risk assessment, FIPS PUB 65 succeeds very well. Its major weakness is that it does not provide detailed guidance on how to actually perform one. It relies in large measure on the expertise and judgement of the members of the risk analysis team, and it permits them wide latitude in selecting the level of detail, structure, schedule and level of effort. Such wide

discretion may lead to significant disparities in the conduct of Risk Assessments in the Federal Government. Using the FIPS 65 methodology, good results require a highly skilled group of ADP experts working together over a considerable period of time. This expenditure of effort may be more than some Federal ADP managers feel comfortable with.

1.2.2 Air Force Risk Analysis Management Program (RAMP)

This document comprises a major part of the Air Force ADP Security Program. The RAMP is still provisional, and undergoing some modifications. It consists of three volumes. The first is a Risk Assessment to be applied to Air Force ADP facilities. The second and third contain the procedures which must be gone through to certify application software (Vol. II) or systems (Vol. III) as appropriate for processing of sensitive data. Thus, for purposes of this comparison, only Volume I is of interest. In a larger context, however, RAMP Volumes II and III are extremely useful, and provide excellent guidance on how to meet military and intelligence certification criteria for processing of sensitive data. The distinction between the contents of Volume I, on the one hand, and Volumes II and III, on the other, is: Volume I evaluates an in-place, operational facility for risks, quantifies the expected loss, and recommends cost-effective countermeasures (i.e., conducts a risk assessment); Volumes II and III provide extensive, detailed guidance on how to certify application software and tactical or C^2 systems when classified materials will be involved. This is a useful distinction and points up the differences between the concerns of general ADP managers and the concerns of the military and intelligence communities. In one sense, the difference is one of degree, since in both cases threats and vulnerabilities are analyzed and assessed. However, a risk assessment for ADP managers usually proceeds on a much higher level (i.e., with less detail) than the detailed, technical testing and review which is often required of systems which will process military and intelligence data. It is a real question what part a risk assessment plays in the certification

effort, and to what extent the tools of the two procedures do or ought to overlap.

The RAMP is a highly structured sequence of procedures set forth in complete detail. Whenever possible, explicit activities are described with painstaking specificity, including the exact procedures to be followed and the output which will be produced by the activity. The major summary statistic is the ALE (Annual Loss Expectancy), although other measures of the level of risk are computed. As usual, the ALE is taken to be the frequency of occurrence of a threat multiplied by the value of the affected asset. Summing over all threat-asset pairs results in a system-wide ALE. In reviewing the document, the following points are worthy of note.

1. Purpose: "The RAMP for ADPFs has been developed to help Air Force managers recognize the threats to their computer operations and the reasonable cost measures that are needed to either prevent, minimize loss, or recover from undesirable events occurring which would reduce Air Force computer resource availability, integrity, or confidentiality. This management awareness with consequent actions is the primary benefit of such risk analyses. The quantification of risk and security measures should also help commanders and the Air Staff allocate resources to protect Air Force computer resources."

2. Level of effort: It is estimated that the initial implementation of the RAMP should take about 60 days, and 30 days thereafter. A substantial number of persons are involved, including the System Security Officer (SSO), the ADP facility manager, a review team assembled by the Base Commander consisting of knowledgeable ADP users and specialists, and support personnel. A conservative estimate would be 4 to 6 man-months most of which would involve highly skilled ADP-trained personnel.

3. An eight-stage approach. The methodology calls for eight separate activities:

I) Inventory of Assets. In this phase, a complete inventory of all assets is performed. Existing inventory records are compared against actual equipment on-hand, and discrepancies resolved. The procedure calls for physically tagging all assets to ensure that this inventory is complete. Asset categories include Hardware (ADP, non-ADP and facilities), Software, Data files, Personnel and Supplies. Each asset is evaluated only in one impact mode (essentially, Loss). This is in contrast to other schemes in which an asset can be valued in more than one impact mode (Destruction, Disclosure, Denial of Service, etc.). A worksheet is prepared in each asset category on which every asset is listed by name and value. A significant feature of this process is that detailed guidance is provided on how to evaluate such items as data files and loss of personnel. For data files, for example, there is a computational algorithm provided for valuing the file based on the fields within a record and the number of records. Very little is left to the judgement of the person conducting the inventory, so that a high degree of consistency should result from following the procedure. The SSO is responsible for conducting this phase.

II) Threat Identification and Evaluation. A work sheet is prepared for every asset (and asset category) identified in Phase I. On this work sheet, the threats which could affect the asset are listed by name. For each threat, two quantities are estimated: 1) the frequency of occurrence of the threat and 2) the magnitude of the threat. The magnitude rating is unique to the RAMP, and represents the percentage of damage which can be expected if the threat occurs. Extensive guidance is provided for magnitude and frequency ratings for many kinds of threats. These estimates are made <u>without</u>

consideration of existing ADP protective measures (i.e., the estimates are made assuming that no ADP-specific countermeasures are in place). The SSO is to receive the guidance of the Resource Protection Committee (RPC), assembled by the Base Commander for this purpose, in making these estimates. It should be noted that none of the threats proposed for evaluation include the more technical types of attack (i.e., O.S. penetration, electronic surveillance, etc.).

(III) Computation of ALE - Unprotected: the asset values in phase I are multiplied by the magnitude and frequency factors identified and estimated in phase II to obtain, for each asset, an annual loss expectancy (ALE). Since these estimates were made assuming that no ADP safeguards are in place, this is the <u>unprotected</u> ALE. Summing over all assets provides the system-wide ALE. The SSO, supported by clerical personnel, performs this activity.

(IV) Inventory and Evaluation of Exisiting Security Measures: the chief (or manager) of the ADP facility performs this phase. The protective measures and devices of the facility are identified, and their annual cost estimated. Summing over all such protective measures yields the total current investment in ADP protective measures and devices. Next, the frequency and magnitude factors estimated in phase II are re-estimated taking into account the presence of the protective measures. Some guidance is provided on how to obtain the list of protective measures. At least two sessions with the RPC are required, as well as walk-throughs and "brain-storming" sessions.

(V) Computation of ALE - Protected: The chief or manager of the ADP facility performs this phase. It is identical to step III, only using the revised frequency and magnitude estimates obtained from phase IV. Note that the difference between the Protected - ALE and the Unprotected - ALE is the annual savings resulting from

the presence of the currently available countermeasures. Hopefully, this savings exceeds their annual cost.

(VI) Measures of Level of Security: During this phase, conducted by the SSO, the ALE, asset values and values of protective measures and devices are manipulated mathematically to provide measures of the current level of security. The most significant is the annual percentage of loss, which is just the ALE expressed as a percentage of the total assets at the facility. A baseline of 10% or less has been established by the Air Force as an acceptable level of risk.

(VII) Selection of Countermeasures: this phase is a lengthy, complex and tedious evaluation of the cost-effectiveness of proposed countermeasures. The methodology consists of comparing the ALE without the countermeasure against the ALE with the countermeasure, computing the difference, and comparing this against the cost of the countermeasure. An interesting feature is that if the current ALE is less than 10% of the total value of all assets, no countermeasures need be considered. Given the complexity and tedious nature of this phase, there is high incentive to ensure that this 10% level has been met. An interesting algorithm based on marginal utility is presented for prioritizing the countermeasures and assessing them sequentially until an acceptable level of risk (i.e., percent of loss) has been obtained. Also, some guidance is given on choosing the most likely candidates for protective measures and devices. The SSO is responsible for conducting this phase, and he reports his findings to the RPC.

(VIII) Action and Reporting: In this phase, the countermeasures recommended for acquisition by the RAMP are acquired (if possible)

and installed. A summary report is prepared for management review, including copies of the working papers prepared as part of the RAMP. All recommended procedural changes are to be rigidly enforced.

The Air Force RAMP as described above has a number of strengths and weaknesses. The strengths include: a highly structured, methodical approach; extensive guidance on evaluation of assets and estimation of threat frequencies and magnitudes; a carefully conceived mechanism for selection and evaluation of countermeasures; unambiguous assignment of responsibility; and review and approval procedures at several points in the process. The potential weaknesses include: a very high level of effort required to carry out the RAMP including extensive participation by the SSO, facility manager and RPC; failure to allow an asset to be evaluated in more than one mode of impact (i.e., inability to distinquish destruction, disclosure, denial of service, etc.); and insufficient emphasis on such areas as encryption, operating system weaknesses, electronic surveillance, segregation of users and data at varying security levels, and hostile code. While these last issues are dealt with in Volumes II and III which examine certification of systems, they should not be totally ignored in a risk assessment. In addition, there is little if any actual experience available using the RAMP, so that its utility is an open issue.

1.2.3 USDA ADP Security Handbook

The risk assessment methodology described in this document was developed for use by the Department of Agriculture in an effort to meet both internal and external (i.e., Privacy Act) requirements and regulations. It is intended for use 1) as a means of assessing the current security position; 2) in raising overall security awareness; and 3) as a management tool for cost-effective allocation of resources. The methodology is very similar to that of FIPS PUB 65, except that orders of magnitude are not permitted, and an effort has been made to include users in the evaluation process. It consists

of twenty pages (of which ten are devoted to a carefully worked-out example), and can be profitably read in less than an hour. The following features merit special attention.

1. Participation of Users: Recognizing that "the only reason for the existence of an ADP facility lies in its function of service to users", the USDA methodology involves users extensively in the impact evaluation process. In particular, users are required to identify critical system assets and services, assess the sensitivity of data files under their control, specify what additional security measures (if any) are required, and estimate the impact associated with the occurrence of major and minor threats.

2. Organized by Threats. A unique feature of this methodology is that the principal categorizing entity is the threat. Each threat is subclassified as "major" or "minor", and a frequency rating is made. Next, an impact estimate is made in each of four impact areas: destruction, disclosure, modification (fraud) and availability (denial of service). This impact rating need not be connected to specific asset values. However, specific assets or groups of assets (e.g., supplies, equipment, etc.) can be identified if more detail is desired. Users also provide estimates of the impact (cost) to them in each impact area. A total impact is thereby obtained in each impact area. The ALE associated with each threat is found by multiplying the impact estimate by the frequency rating. Summing over all such threats yields the system-wide ALE. Orders of magnitude are not used. Note that system assets are never explicitly inventoried or evaluated.

3. Loosely Structured. As with FIPS 65, the instructions are left sufficiently general so that many decisions concerning level of effort and detail are left to the team conducting the risk assessment.

4. Selection of Countermeasures. A positive feature of this methodology is an orderly approach to the selection of countermeasures. A matrix is prepared listing threats on one axis and countermeasures on the other. An entry in the resulting table indicates that the given countermeasure can be effective in preventing the designated threat. Countermeasures which are effective against many threats (or against the most serious threats) are then good candidates for implementation. As usual, the countermeasure is evaluated by recomputing the ALE with the countermeasure in place, and comparing the resulting savings against the cost of the countermeasure.

The major weaknesses in the USDA methodology is lack of guidance in arriving at a suitable list of threats and in estimating their frequency and impact. A team approach is recommended, but the size and content of the team is left open. While a number of activities are described, the methodology is not a phased approach. Activities may overlap, and responsibility for parameter estimation is split among different groups. Given a strong prior foundation, the countermeasure evaluation and selection mechanism is quite sound and orderly. A major strength of the methodology is the active involvement of users in the evaluation of the impact resulting from threat occurrence. As in FIPS 65, the directions for implementation are sufficiently general that assessments involving widely differing levels of detail and effort could reasonably be said to fit under the same description. While this allows ADP managers flexibility in choosing the amount of time and resources to be allocated, such decisions will have a great impact on the validity and utility of the results of the assessment.

1.2.4 AFIPS Checklist for Computer Center Self-Audits

This document should be required reading for anyone concerned with issues of computer security. It consists of 954 "embarassing questions" relating to all facets of ADP operations. These are broken down into 9 major sections: Planning and Risk Analysis; Physical Security; Backup and Recovery;

Administration Controls; Systems Hardware and Software; Communications; Distributed Risk; Applications; and Security Audit. Each section is introduced by several pages of discussion of terminology, concepts, available technology and recommended practices. Then follows a detailed checklist forcing the reader to consider a multitude of disturbing and unlooked-for possibilities. To read the book carefully is to see your system through the eyes of your enemy - the strengths, weaknesses, targets and avenues of attack are brought into sharp focus. In addition, the dangers from carelessness, errors, sloppy procedures and poor planning are highlighted. It truly describes an ADP manager's nightmare!

Insofar as a major purpose for conducting a risk assessment is to provide a careful, complete review of the security properties of a facility and to heighten security awareness, the AFIPS manual provides an extremely effective and complete mechanism to accomplish this. However, it does not provide the quantitative measure of degree of risk (i.e., ALE) which results from a formal Risk Assessment. It is not a quantitative tool; it is qualitative in its purposes and mechanisms. Its value lies in the completeness of its treatment of the security posture of a facility or system -- the confidence that, having worked through the checklist, no significant security feature will have been left unexamined. As such, it is a useful reference while conducting a quantitative Risk Assessment. It may be compared to a Risk Assessment Methodology insofar as that Methodology does or does not permit and encourage the same degree of completeness of review. This is particularly true when those conducting the risk assessment are new to the issues of computer security or are not ADP professionals.

1.2.5 Navy RAM

This document contains a Risk Assessment Methodology (RAM) which was prepared by SDC for the Department of the Navy. It has been proposed for adoption and use by Navy personnel in conducting Risk Assessments of Navy

ADP systems and facilities. The unique mission and requirements of the Navy influenced the development of the methodology in several ways. On the one hand, it was desired that the methodology be useable by Navy personnel who were not highly skilled in the area of computer security. On the other hand, adherence to sound technical and theoretical principles are of central importance, including reliability of results and use of quantitative measuring techniques. Thus, detailed descriptions of RAM activities were required as well as extensive guidance in identifying and evaluating system threats and vulnerabilities.

In order to accomplish these objectives, a number of innovative techniques were devised for use in the Navy RAM. Probably the most significant of these was to separate out the vulnerabilities of the system or facility as a separate entity for evaluation. The vulnerabilities are seen as paths or openings by means of which a threat can reach a target asset to cause loss or damage. A low vulnerability rating indicates that only a few attacks can succeed in exploiting the vulnerability. Likewise, a high rating indicates that most attacks against that vulnerability will be successful. Thus, mathematically, the vulnerability rating functions similarly to the "Magnitude" factor in the Air Force RAMP. It reflects the fact that strong safeguards will succeed in deflecting a large percentage of attacks, thereby reducing the ALE.

The Navy RAM consists of six phases. In the initial three phases, the Threats, Vulnerabilities and Assets of the system are identified and evaluated. Extensive guidance is provided in making these evaluations, including checklists, pointers to sources of information, scenario construction and explicit estimation guidelines. A generic list of Threats and Vulnerabilities has been compiled, although users may add others as appropriate. Assets are inventoried, and are evaluated in each of four modes of impact: Destruction, Unauthorized Disclosure, Modification and Denial of Service. In all cases,

orders of magnitude are used. Vulnerabilities are rated using qualitative verbal descriptions (VL, L, M, H, VH), and are incorporated into the algorithm via a look-up table. Because the impact of certain threats is extremely difficult to evaluate in dollars (e.g., disclosure of classified data), the option is given of evaluating assets using a qualitative, non-dollar-valued technique. This feature is optional, and is used only at the discretion of the person conducting the risk assessment.

The fourth phase involves matching Threats against Vulnerabilities to obtain plausible Threat/Vulnerability attack scenarios. This is done in each of the four modes of impact. The Threat/Vulnerability merger forms which accomplish this have been prepared for the generic categories supplied with the RAM. The user who requires additional entries must also modify these forms accordingly. At this stage, the vulnerability rating is matched against the threat frequency rating to obtain a successful attack frequency. This is done via a look-up table. As indicated, low vulnerability ratings will considerably decrease the attack frequency; high ratings will leave it largely unchanged.

In the fifth stage, the attack scenarios obtained in phase four are matched against target assets. As in all methodologies, the frequency of attack multiplied by the impact (in this case, asset value) yields the ALE (Annual Loss Expectancy). Summing over all T-V-A triples yields the system-wide ALE. In the Navy RAM, eight forms are provided for this purpose: four for each of the impact modes, and again broken down by dollar vs. non-dollar evaluation of assets. For the non-dollar-valued assets, the level of risk is obtained by summing the successful attack frequencies of all T-V pairs which affect the asset, obtaining in this way a total successful attack frequency. Management can then decide whether that level of risk is acceptable. It should be noted that all computations are table-driven, and use orders of magnitude. Also, the matching of T-V pairs against target assets is left to the judgement of the person conducting the Risk Assessment. This

can be a lengthy and tedious undertaking. Finally, the forms permit useful intermediate results, such as computing the ALE associated with a specific threat or vulnerability. This can be useful in analyzing the overall security posture of a facility and in making educated guesses concerning areas where additional protective measures are needed.

The sixth stage consists of selection of countermeasures. As in other risk assessments, the technique is to recompute the ALE assuming the countermeasure is in place, and then to compare the resultant savings with the cost of the countermeasure. In the Navy RAM, the countermeasure enters into the algorithm by reducing the vulnerability rating (and thereby, the successful attack frequency). A reference manual on potential countermeasures has been prepared to aid in this process.

One final feature of this methodology may be used. When threat frequencies and asset values are estimated, a precision rating is called for. The precision rating indicates how accurate the evaluator feels his estimate to be. By using this precision rating, a new ALE can be computed on the assumption that the evaluator consistently estimated below the true ratings. The resultant ALE will therefore be correspondingly higher. This process is termed the "Worst Case Analysis", and the RAM warns that this result should be viewed with a jaundiced eye. It may be thought of as the RAM's version of Murphy's Law. This second estimate is useful insofar as it gives an indication of the range of reasonable values to be considered.

2. COMPARISON OF METHODOLOGIES

In this section, the methodologies will be compared in each of five major comparison criteria.

2.1 SOUNDNESS OF UNDERLYING THEORY

There are three aspects which must be examined under this heading. They are completeness (do the parameters identified by the methodology provide a complete profile of the security and integrity posture of the facility or system being assessed?), accuracy of estimation (are the mechanisms used by the methodology to estimate parameters deserving of confidence?), and validity (is the computational algorithm used to combine parameter estimates into summary statistics a reasonable model of reality?). We consider each in turn.

2.1.1 Completeness

Each of the methodologies calls for the identification (as opposed to subsequent estimation) of lists of objects, events, situations, and relationships. Subsequently, numerical estimates (of values, frequency, impact, etc.) may be associated with each of the members of these lists for computational purposes. However, if the lists themselves are incomplete, or fail to reflect a significant aspect of the security and integrity posture of the facility or system, then subsequent results will be, to that extent, in error. Hence, the mechanisms provided by the risk assessment for construction of such lists, and any features which help to assure completeness, are of concern.

The FIPS 65 methodology calls for the preparation of lists in the areas of assets (for replacement costs), threats, protective measures, and asset-threat pairs (for each asset, listing of the threats which could affect it). With the exception of threats, for which an extensive list of examples is provided in the Appendix, the methodology provides no explicit features for assuring that the resulting inventories are comprehensive. For example,

a crucial step is the matching of assets against potential threats. The document says: "All of the organizations applications systems, or data files arranged by application, should be listed on the worksheet(s). By tracing the flow of data through a system, the team will be able to pinpoint where in the processing the threats identified in the preliminary study could occur." This is followed by three pages of general guidance on how to estimate parameters, but the initial step of identifying the threat-asset pairs is never mentioned again. Clearly, the methodology relies almost entirely on the expertise of the risk assessment team to guarantee that the pairings have been done in a reasonable and comprehensive manner.

The Air Force RAMP is extremely strong in this area. Pre-printed forms are required in each of the areas of asset identification and evaluation, threat identification, threat-asset pairing and inventory of countermeasures. In addition, explicit lists and activity descriptions are provided indicating unambiguously and in detail exactly the procedures to be followed in compiling these lists. Finally, intermediate products are reviewed by a highly qualified, skilled committee (the RPC) for completeness, reasonableness and consistency. This is a wise use of manpower since the committee is only assembled at intervals to review drafts and aid in the more difficult decisions; the bulk of the work is conducted by the SSO and his assistants. Thus, while the RAMP does not provide detailed, extensive checklists of ADP threats and vulnerabilities, it does provide a formal review mechanism and specificity of procedures as the means of assuring completeness. This is a reasonable approach to the problem.

The Risk Assessment developed by the Department of Agriculture only requires two lists: 1) threats to the facility, and 2) potential countermeasures. The assets of the facility, are never explicitly inventoried or evaluated. Very little guidance is provided on how to obtain completeness in these lists. Some examples are given, but there is great reliance on the skill

and judgement of those individuals conducting the risk assessment. The methodology does not explicitly inventory or even discuss vulnerabilities in ADP systems or facilities.

While this issue is not strictly applicable to the AFIPS checklist, systematic use of the material in this document would provide very high confidence that every important facet of ADP security and integrity had been considered. An explicit method or procedure for incorporating this checklist into less structured risk assessment methodologies would be a major improvement.

The approach taken to this problem by the Navy RAM is to incorporate pre-established lists explicitly as part of the methodology. While the individuals conducting the risk assessment can add to these lists if they choose, they need not do so. The lists so prepared and utilized are comprehensive, and deal with the full spectrum of general threats and vulnerabilities in ADP systems and facilities. There are, however, three places where the judgement and expertise of the individual is called for in the preparation of lists: 1) in preparing the inventory (and evaluation) of system assets; 2) in identifying which assets are impacted by a given threat/vulnerability attack pair; and 3) in preparing a list of potential countermeasures. While extensive guidance is provided for 1) and 3), 2) is a lengthy and potentially frustrating activity for which only general, non-specific rules and advice can be given.

2.1.2 Accuracy of Estimation

Each of the methodologies (except the AFIPS checklist) is at least partially concerned with quantifying the risk to an ADP system or facility. As such, each requires that certain parameters (previously identified) be numerically estimated. These estimates will subsequently enter into certain comparisons and arithmetic manipulations. Hence, the accuracy with which these estimates can be made is a major concern, and the methodologies must be examined concerning the support and guidance they provided to users in this area.

Before turning to the actual methodology documents, however, a general comment should be made. All of the methodologies recognize explicitly that the process of estimating such things as frequency of occurrence (say, of theft) or degree of impact (say, of disclosure of personal data) is necessarily and unavoidably inexact. One method which has been proposed to help those conducting risk assessments is the use of orders of magnitude. The use of such a device speeds the estimation process and (presumably) uses manpower more effectively. However, for parameters which can be estimated very precisely (as, for example, the replacement cost of a disk drive), such accuracy is irretrievably lost when orders of magnitude are employed. In general, use of orders of magnitude, at least on a first go-around, seems to make sense given the highly imprecise nature of many of the estimates to be made. This technique is employed both by FIPS 65 and by the Navy RAM.

The most significant means by which a methodology can aid in assuring the accuracy and reasonableness of parameter estimates is to provide explicit guidance on making or obtaining these estimates.

The FIPS 65 methodology requires the risk assessment team to estimate the frequency and impact of every threat affecting a given asset in each of the four impact modes (destruction, modification, disclosure and denial of service). That is, four frequency and impact estimates are made for each threat-asset pair. No specific guidance is given about how these estimates are to be made. The only aids are in the nature of advice, encouragement and mention of some common pit-falls. A similar situation exists when the risk assessment team must estimate the reduced ALE when a particular countermeasure is in place. The methodology relies almost entirely on the experience and expertise of the members of the Risk Assessment team for the validity of the parameter estimates.

The Air Force RAMP requires the estimation of the value of each asset, the frequency of occurrence of each threat, and the magnitude (percentage of assets destroyed) of each threat. The last two quantities are subsequently re-estimated assuming various countermeasure configurations. The RAMP provides extensive guidance on making these estimates. The most remarkable example is the detailed algorithm for evaluating the sensitivity of data files. However, frequency and magnitude ratings for earthquakes, storms, tornadoes, lightning, etc. are furnished in the appendices. The RAMP also provides an interim review of the estimate by the RPC for reasonableness and consistency. Orders of magnitude are not used. Finally, the RAMP provides specific references to potential sources of information not already provided, and explicitly requires those conducting the assessment to contact these sources and justify their ratings.

The methodology developed by the Department of Agriculture is similar to FIPS 65 in that it calls for the estimation of frequency of occurrence of threats, and the estimation of impact in each of four impact categories. A unique feature of this methodology, however, is that users are explicitly involved in the estimation of impact. Thus, the results of the assessment reflect not only operating costs, but also the costs to those who use the facility and rely on its processing. This approach makes particularly good sense when the users will also bear part (or all) of the cost of installing additional countermeasures. The methodology does not give any explicit guidance on how to make these estimates, relying entirely on the experience and judgement of those conducting the assessment. Orders of magnitude are not used.

The issue is not applicable to the AFIPS checklist.

The Navy RAM requires estimation of: 1) frequency of occurrence for Threats; 2) level of vulnerability for Vulnerabilities; 3) values in each of four

impact modes for assets; and 4) effectiveness ratings for countermeasures. Orders of magnitude are permitted in 1) and 3), and qualitative verbal descriptors are used in 2) and 4). In addition, assets for which the assignment of a dollar value is particularly difficult or misleading are allowed to be non-dollar valued. The Threat and Vulnerability forms contain extensive, detailed guidance on how to arrive at a reasonable, accurate estimate in each separate case. Only general guidance is provided for evaluation of assets. A manual containing descriptions and ratings of countermeasures is also provided. Finally, estimates of threat frequencies and asset values can be accompanied by a precision rating indicating the confidence the estimator has in his estimate. This can be very useful in the analysis of results. In general, the methodology has gone to great lengths to aid implementors in arriving at reasonable, accurate estimates of the critical parameters.

2.1.3 Algorithmic Validity

Once the critical parameters are identified and estimated, they are combined into summary statistics using some computational algorithm, however simple or complex it may be. The results, therefore, are only as meaningful as the underlying mathematical model. On the one hand, simplicity and intuitive appeal are strengths, since they lend confidence to a reader or user that he understands what the figures mean and how they were arrived at. On the other hand, the methodology should have a structure which is sufficiently rich that the security features of the system can be modeled accurately and reliably. Finally, any computational aids which minimize the mathematical burden on the user must also be considered to be strengths.

The FIPS 65 methodology has an eminently reasonable, structured approach.

[Frequency of occurrence] X [Impact] = [Annual Loss Exposure].

The methodology, as we have seen, uses orders of magnitude to estimate frequencies and impact, and provides a look-up table to perform the required multiplication. Note that, for each asset, the threats which can affect it are listed. This essentially identifies a large number of asset-threat pairs. For each such pair, frequency and impact ratings are made in each of the four modes of impact. This means that an extremely large number of ratings are made. In most of the methodologies, a frequency rating is made once for each threat, and then reused as the threat is matched against each asset. Such a time-saving device is not a feature of the FIPS 65 methodology. While re-estimation for each case is probably more accurate from a theoretical point of view, it imposes a sizeable burden on those conducting the risk assessment.

The next step is to sum in various directions. Summing over the four impact modes yields an ALE for each threat-asset pair; summing over all threats yields total ALE for each individual asset; summing the contribution of a given threat over each of the assets it affects yields the ALE associated with a given threat; and summing over all threat-asset pairs yields the systemwide total ALE. This procedure has a strong, intuitive appeal. Its principal weakness lies in the extremely large number of estimates which must be made. In addition, the parameters which are required are slippery, even to ADP experts, and the methodology gives little guidance in identifying and estimating them. Finally, the fact that estimates are to reflect frequencies and impact assuming none of the current protective measures is somewhat strange, and will tend to complicate the process unnecessarily.

The Air Force RAMP has two theoretical features which merit special examination. First, the RAMP only permits a given asset to be valued and impacted in one impact mode. For example, equipment is valued by replacement cost, data by sensitivity (i.e., disclosure), personnel by level of training and experience, etc. Nevertheless, any given asset is only evaluated in one mode, so that, for example, the impact of denial of service cannot be separated out from the

impact of unauthorized disclosure, or again from loss of data integrity (modification). While the text of the RAMP recognizes these alternative events, no structure is provided in the RAMP itself to model or evaluate them. This is a serious weakness.

Second, the RAMP requires the insertion of a factor called the magnitude of a threat. This magnitude is used to represent the percentage of an affected asset which would be damaged or destroyed, and is estimated once for each threat, regardless of which assets are being affected. While it has a certain intuitive appeal, and makes sense for some threats, the percentage of assets being affected is often more a characteristic of the facility than it is of the threat (e.g., a brick building will not be as extensively damaged by a fire as a wooden building would be). On the other hand, the attempt to model the difference between "severe" and "mild" threats is valid. The Department of Agriculture methodology accomplishes this by distinguishing two threat categories - "major" and "minor" - and recomputing frequency and impact for each category. FIPS 65, of course, recomputes frequency and impact for every threat-asset pair, so that the problem does not arise. The decision whether to use a magnitude rating, and which threats most reasonably are modified in this way, is a matter both of policy and practical experience.

As expected, the Air Force RAMP matches potential threats against likely target assets, obtaining in this way a large number of threat-asset pairs. Asset value multiplied by threat frequency and the magnitude factor yields the ALE for each threat-asset pair. The frequency rating is re-used if a threat affects more than one asset, so that the total number of estimates is dramatically reduced vis-a-vis FIPS 65. Total ALEs are obtained by summing first over threats, and then over assets, to obtain a system-wide ALE. As in FIPS 65, initial estimates are made assuming no countermeasures (even those currently in place are not considered in the initial phase). Estimates are subsequently revised to reflect (first) current and (second) proposed countermeasures, and the ALE recomputed in each case. Finally, the

RAMP contains an elaborate and tedious countermeasure selection mechanism based on marginal utility metrics (analysis of benefit-to-cost ratios resulting from incremental countermeasure additions). This algorithm would seem to be unjustified in light of the theoretical weaknesses pointed out above. A useful feature is the baseline risk assumption (10% of total assets). Many facilities may find such a concept meaningful, and it can save the risk assessment team much unnecessary labor. No explanation is provided to justify the 10% figure.

An interesting feature of the Department of Agriculture methodology is that the fundamental organizing unit is the Threat. The frequency of each Threat (major and/or minor) is estimated once, and the impact of the threat is estimated in each of four impact modes. Note that affected assets are not required to be explicitly identified or evaluated, although the assessors may do so if they choose. It is very difficult to decide exactly where these estimates come from, since the level of detail is so coarse. While this method of organization facilitates (indeed, requires) that partial ALEs be computed for each threat, partial ALEs by asset are not available. In general, the algorithm appears to be appropriate to the level of detail required by the rest of the methodology. This is a high-level approach, and great accuracy is neither expected nor required. The actual number of estimates made in this methodology is least of all those examined: one frequency and four impacts for each threat.

This criteria is not applicable to the AFIPS checklist.

The Navy RAM has devised an algorithm which is intended to provide an extremely rich level of structure for modeling attack scenarios while, at the same time, holding down the actual number of estimates to a minimum. As discussed, the methodology identifies a very large number of threat-vulnerability-asset (T-V-A) scenarios, and this in each of the four modes of impact. The threat frequency modified by the vulnerability level and

multiplied by the asset value yields the ALE associated with the given T-V-A triple. Summing along a fixed axis (by threat, by vulnerability, or by asset) yields useful partial ALEs.

Two important points need to be made about this algorithm. First, while it is true that T-V-A scenario triples are identified separately for each of the four impact modes, the <u>same</u> threat and vulnerability ratings are used no matter which mode is under consideration (assets are valued separately for each mode). Thus, it is assumed by the model that, for example, power outages resulting in destruction of equipment will occur with the same frequency as power outages resulting in denial of service. While this assumption saves effort (since frequencies and vulnerability levels do not have to be re-estimated in each impact mode), it probably distorts the final results to some extent, and is thus a theoretical weakness.

Second, the full value of the asset is used whenever it appears in a T-V-A scenario. That is, the measure of "impact" used by the Navy RAM is the asset value. This is similar to the approach taken by the Air Force RAMP, where the asset value entered explicitly into the algorithm (indeed, was taken to be synonomous with impact). Note that, in the RAMP, each threat-asset pair is modified by the presence of a magnitude factor indicating the severity of the threat. In the Navy RAM, however, it is assumed that every threat is "totally severe", i.e., results in total destruction of or damage to any potential asset. This is probably an unrealistically pessimistic view of the situation, even allowing that a conservative estimate is preferrable to an overly optimistic one. Note also that in the FIPS 65 and Department of Agriculture methodologies, asset values are not required and do not explicitly enter into the "impact" estimate as they do in the Navy RAM and Air Force RAMP. "Impact" is simply defined as "cost associated with the occurrence of a threat", and is thus only loosely connected to asset values.

The use of orders of magnitude, qualitative verbal descriptions, look-up tables for performing computations and easy-to-use pre-printed forms makes the actual computation using the Navy RAM extremely simple. Such a table-driven approach is especially valuable when non-skilled personnel are assigned to carry out the risk assessment.

2.2 LEVEL OF MANPOWER REQUIRED

There are two related issues here. The first concerns the level of expertise of those who will conduct the assessment. The second concerns the total amount of time required to conduct the assessment.

One difficulty in evaluating this category for FIPS 65 and for the DoA Handbook is that both documents are quite vague on this point -- perhaps deliberately so. The Department of Agriculture methodology is never more precise than "facility personnel" and "users". The FIPS 65 methodology speaks about "highly skilled manpower" and gives a suggested list of organizational components to be represented on the team. If FIPS 65 is to be taken at face value, the risk assessment team will be drawn from extremely knowledgeable and responsible ADP professionals. The costs associated with the Risk Assessment will be correspondingly high.

Another indication that both FIPS 65 and the DoA Handbook require highly skilled personnel and substantial amounts of time is that neither of them give detailed guidance on how to construct the inventories or estimate the critical parameters. That is, both methodologies rely very heavily on the expertise and judgement of the risk assessment team. Thus, a reasonable expectation is that the team merits this trust, both in terms of the professional accomplishments of its members and in terms of the time and money given it to complete its task.

The Air Force RAMP is quite explicit about assigning specific duties and estimating the amount of time involved. The SSO and the ADP facility manager share principal responsibility, assisted by support personnel as required, and reviewed at frequent intervals by the Resource Protection Committee (RPC). The RPC consists of high-level, responsible managers and officers. The RAMP estimates an initial period of 60 days to conduct the Risk Assessment. Subsequent assessments should only require 30 days. A reasonable estimate would be 4 to 6 man-months to conduct the initial assessment.

The AFIPS checklist can be read in about three days, and could be applied to a facility by a knowledgeable person in about a week. If this individual were in a responsible position, the weaknesses which were detected could be corrected almost immediately. However, budgetary justification of additional expenditures would require a quantitative measure *not provided* by the AFIPS checklist. In this regard, it would be extremely interesting to see the Risk Assessment, itself, cost-justified using its own techniques. It is quite possible that informal methods such as the AFIPS checklist are more cost effective!

A design goal of the Navy RAM is that it be useable by lower-level personnel who are not necessarily trained in ADP security methods or technology. The assessor is led step by step through a sequence of discrete activities none of which require special training or expertise. Thus, while the methodology takes certain liberties in terms of abstract, theoretical correctness, it provides a remarkable quality of result given unskilled personnel and a short period of time. In addition, where more time and money are available, the methodology provides a way of structuring the approach into manageable pieces which can be examined individually, and later combined to obtain quantitative measures of risk. Given the goal of obtaining an acceptably accurate estimate of risk without great expenditure of time and expertise, the Navy RAM is clearly superior to any of the other methodologies examined.

2.3 EASE OF USE, COMPLETENESS OF INSTRUCTION AND DESIGNATION OF RESPONSIBILITY

To appreciate the significance of this category, consider the following three examples giving "instructions" on threat identification and estimation.

Example A. Identify the threats to the facility and estimate their frequency of occurrence.

Example B. From the following list, select those threats which impact the facility, and use the accompanying general guidelines and criteria to estimate the frequency of occurrence of each threat.

Example C. Each of the following threats can impact an ADP facility. Review them, and follow the detailed instructions accompanying each one on how to estimate their frequency of occurrence for the particular facility in question.

It becomes clear from this that there is a wide spectrum of possible ways of "giving an instruction". In general, the FIPS 65 and Department of Agriculture methodologies most nearly resemble A and B. On the other hand, the Air Force RAMP and the Navy RAM most nearly resemble B and C.

In general, a methodology which follows most closely to Example C above will be easier to use in the sense that fewer decisions are left to the individual conducting the assessment. He has a well-defined sequence of simple activities to perform which are structured in such a way that the final result is the ALE being sought. In addition, the activities he must perform have been mapped out completely in advance, together with the results which must be produced and evaluation criteria. In this sense, methodologies which adhere more closely to Example C will hold their assessors more accountable for the results of their efforts. In the Navy RAM, one part of each estimation procedure is to justify the rating based on the specific criteria which

have been set forth. In the Air Force RAMP, the risk assessors are subject to review and scrutiny of intermediate results at every stage.

In comparison, the instructions given in FIPS 65 and the Department of Agriculture methodologies are "high level". Most of the actual decisions concerning level of detail, methods of organization, accountability and methods of presentation are left to the judgement and discretion of the Risk Assessors. As a result, they are apt to have to expend a considerable amount of time and effort deciding <u>how</u> to do the assessment rather than actually <u>doing</u> it. In addition, because the guidelines are so general and vague, it is hard to argue that whatever is produced is insufficient. That is, so much can come under the general instructions as presented that very little is excluded. In this sense, a manager is fortunate indeed if the result of an assessment using either of these methodologies meets his needs.

2.4 USE AND RELIABILITY OF QUANTITATIVE METRICS

The fundamental measure of level of risk is generally agreed to be the Annual Loss Expectancy (ALE). With the exception of the AFIPS Checklist, the other methodologies reviewed did, indeed, manage to compute this quantity one way or another. Along the way, however, some of the methodologies computed other quantities which are also useful, and which can also indicate the level of risk faced at an ADP system or facility.

- Frequency of Attack: In FIPS 65, this is re-estimated for each impact mode <u>and</u> against every asset. The number of estimates is thus

  (number of threat-asset pairs) X 4.

  In the Air Force RAMP, DoA Handbook and Navy RAM, the threat frequency is rated <u>once</u> and reused for each of the various assets and impact modes where the threat appears.

- Magnitude of Threat: This quantity is unique to the Air Force RAMP. It is estimated explicitly for each threat, and indicates the "severity" of the attack. This quantity is subsumed in the "impact" estimate in FIPS 65. In the DoA methodology, it appears as the "major" and "minor" threat categories. It is completely missing in the Navy RAM.

- Impact of Threat: In FIPS 65, this is re-estimated for every threat-asset pair, and for each of the four impact modes. In the DoA methodology, it is estimated for each threat, period; there is no algorithm for its computation. In the Air Force RAMP, it is simply the (single) asset value. In the Navy RAM, it is likewise the asset value, but a separate value is estimated in each of the four modes of impact.

- Vulnerability Level: This quantity is unique to the Navy RAM, and is an explicit rating of the degree of resistance of a facility to attack along a number of potential paths (Vulnerabilities).

- Frequency of Successful Attack: In all methodologies except the Navy RAM, this is the same as the frequency of attack. In the Navy RAM, it is the attack frequency modified by the Vulnerability Level.

- Individual ALE's: In FIPS 65, every threat-asset pair in each of the four impact modes has an ALE. It is simply the impact estimate multiplied by the frequency estimate. In the DoA methodology, ALE is computed for each threat. It is the product of the frequency of occurrence of the threat multiplied by the impact of the threat. A separate impact, and hence, ALE, is computed for each of the four modes of impact.

In the Air Force RAMP, the ALE is computed for each threat-asset pair. It is the product of the frequency, magnitude and asset value estimates. In the Navy RAM, one ALE is computed for each T-V-A scenario triple. It is the product of the attack frequency (modified by the vulnerability level) and the asset value. A separate ALE is computed in each impact mode.

- Unprotected ALE's: In FIPS 65 and the Air Force RAMP, the ALE is initially computed assuming that no currently installed countermeasures are in place. The ALE is then recomputed taking these measures into account. In the other methodologies, current, in-place countermeasures are not ignored on the first go-around.

- Level of Vulnerability: In the Navy RAM, a mechanism is provided to compute the number of successful attacks against sensitive non-dollar-valued assets.

- Economic Measures: Each of the methodologies provides for a mechanism to evaluate whether proposed countermeasures are cost effective. In all cases, this measure amounts to comparing the (pro-rated) cost of the countermeasure against the savings it would bring about. The savings are found by recomputing the ALE assuming the countermeasure is in place, and subtracting this figure from the original ALE. The Air Force RAMP also considers the marginal utility of each countermeasure, and uses this figure to prioritize them.

- Precision Ratings: The Navy RAM permits a precision rating to be associated with each frequency and value estimate. These are subsequently available to use in the Worst Case Analysis.

An important consideration as these "metrics" are reviewed is to what extent a methodology will be consistent in computing them. That is, imagine two different independent but equally expert teams evaluating the same facility: is it reasonable to expect their results (i.e., the measures they finally compute) will be close to each other. The answer will depend in large part on how explicit the instructions and guidelines are. The more decisions and judgements which are left to the discretion of the team, the farther apart the resulting estimates are likely to be. The issues discussed in 2.3 are apropos here, and it appears that results obtained from the Air Force RAMP and the Navy RAM are much more likely to be consistent and reliably comparable than are the results obtained using the other two.

2.5 APPROPRIATENESS FOR USE BY THE DEPARTMENT OF NAVY

One of the methodologies under consideration (the Navy RAM) was actually written explicitly for and under contract to the Navy. It is peppered with examples and terminology specific to Navy and military ADP operating concerns. Classified materials, denial of service, special ship-board concerns and military ADP requirements were all explicitly taken into account in preparing the lists of generic threats and vulnerabilities and in providing estimation guidance. Thus, the Navy RAM as it appears in the Handbook is in truth a Navy document reflecting the unique mission and ADP requirements of the Navy.

Of the remaining three methodologies (excluding the AFIPS checklist), the one most nearly suited to Navy use is the Air Force RAMP. It also is highly structured and procedural, and it also reflects many uniquely military concerns (e.g., classified data). It requires substantially more effort to perform than the Navy RAM, but could probably be tailored for Navy use with relatively little effort.

Neither of the remaining methodologies (FIPS 65 and the DoA Handbook) are suited for Navy use. While they are sufficiently general that they could be used on almost anything, they do not provide the low-level, detailed guidance on frequency and impact evaluation which is needed in a military setting.

3. CONCLUSIONS

The analysis shows that the Navy RAM compares very favorably to the other methodologies encompassed in this comparison. It is highly structured and procedural; it contains explicit guidelines to aid users in estimating critical parameters; it explicitly sets forth system vulnerabilities for examination and evaluation; it does not require a great deal of time or ADP expertise to conduct; it uses orders of magnitude and verbal qualitative descriptors to simplify the user's task; all computations are performed by look-up tables; it adheres closely to the quantitative metrics needed by management for selection of cost-effective countermeasures; and it provides useful special features such as non-dollar-valued assets and a worst case analysis technique.

The major competitor of the Navy RAM was the Air Force RAMP. Like the RAM, it is highly structural and procedural, and it too offers excellent guidance on threat frequency ratings. A potentially useful feature is the use of a "magnitude factor" as a measure of the severity of impact of a threat. However, this is more than offset by the failure to distinguish alternative modes of impact, and the insistence on equating impact with asset replacement value. Finally, the RAMP explicitly requires a very high expenditure of skilled manpower, including frequent review of all intermediate results by a supervising committee. It is arguable that such a level of effort is not required given the unavoidably imprecise nature of a risk assessment. However, the participation of management does increase the likelihood that the recommendations of the RAMP will be acted upon.

The two civilian methodologies, FIPS 65 and the DoA Handbook, both suffer from a number of flaws. Principal among these is the fact that both methodologies are very loosely structured and overly general in their descriptions of tasks. They rely to a great extent on the judgement and expertise of those conducting the risk assessment. They offer very little guidance for the identification and evaluation of critical parameters. While FIPS 65

uses orders of magnitude to ease the task of estimation, it requires estimation of an enormous number of parameters (an attack frequency and impact estimate for every threat-asset pair in each impact mode). Thus, it does not attempt to take advantage of any simplifying techniques to ease the burden on the user. By contrast, the DoA Handbook does not explicitly use assets at all in its algorithm. Its major strength is that it does explicitly involve users in the estimation of costs associated with threat occurrence, and that it presents an exceptionally clean and appealing way of organizing the countermeasure selection process.

Both the FIPS 65 methodology and the Air Force RAMP require initial computation of the ALE in the absence of existing and currently installed countermeasures. There appears to be no benefit sufficiently great to justify this process.

Finally, the AFIPS checklist is strongly recommended as an excellent means of obtaining awareness of the issues and techniques of ADP security. To the extent that a risk assessment is used to heighten security knowledge and awareness, the AFIPS checklist could prove to be extremely useful. While it cannot replace a risk assessment (since it is non-quantitative in nature and technique), it could certainly supplement one, and it provides a risk assessor with an excellent source of potential threats and vulnerabilities.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: TM-WD-7999/202/00

2. GOVT ACCESSION NO.: AD-A089346

3. RECIPIENT'S CATALOG NUMBER:

4. TITLE (and Subtitle): A COMPARISON OF RISK ASSESSMENT METHODOLOGIES

5. TYPE OF REPORT & PERIOD COVERED: FINAL rept.

6. PERFORMING ORG. REPORT NUMBER: TM-WD-7999/202/00

7. AUTHOR(s): H. Stephen Morse

8. CONTRACT OR GRANT NUMBER(s): N00173-78-C-0455

9. PERFORMING ORGANIZATION NAME AND ADDRESS:
System Development Corporation
7929 Westpark Drive
McLean, Virginia 22102

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: Task #8

11. CONTROLLING OFFICE NAME AND ADDRESS:
Naval Research Laboratory
Washington, D.C. 20375

12. REPORT DATE: 30 June [redacted]980

13. NUMBER OF PAGES: 42

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): 45

15. SECURITY CLASS. (of this report): Unclassified

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE:

16. DISTRIBUTION STATEMENT (of this Report):

Approved for public release.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report):

18. SUPPLEMENTARY NOTES:

19. KEY WORDS (Continue on reverse side if necessary and identify by block number):

Risk Assessment, Risk Analysis, Methodology Comparison, Risk Management

20. ABSTRACT (Continue on reverse side if necessary and identify by block number):

This document contains a comparison of five Risk Assessment Methdologies. They are: FIPS 65; the Air Force RAMP; Chapter 6 of the DoA ADP Handbook; the AFIPS Self-Audit Checklist; and the RAM developed by SDC for the Navy. These methodologies are compared on a number of criteria including theoretical soundness, level of effort required, ease of use, completeness of instructions, use and reliabil[illegible] of quantitative metrics, and appropriateness for use by the Department of the Navy.

DD FORM 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified 339860